

FUNCTORIAL HIERARCHICAL CLUSTERING WITH OVERLAPS

JARED CULBERTSON, DAN P. GURALNIK, AND PETER F. STILLER

ABSTRACT. This work draws its inspiration from three important sources of research on dissimilarity-based clustering and intertwines those three threads into a consistent principled functorial theory of clustering. Those three are the overlapping clustering of Jardine and Sibson, the functorial approach of Carlsson and Mémoli to partition-based clustering, and the Isbell/Dress school's study of injective envelopes. Carlsson and Mémoli introduce the idea of viewing clustering methods as functors from a category of metric spaces to a category of clusters, with functoriality subsuming many desirable properties. Our first series of results extends their theory of functorial clustering schemes to methods that allow overlapping clusters in the spirit of Jardine and Sibson. This obviates some of the unpleasant effects of chaining that occur, for example with single-linkage clustering. We prove an equivalence between these general overlapping clustering functors and projections of weight spaces to what we term clustering domains, by focusing on the order structure determined by the morphisms. As a specific application of this machinery, we are able to prove that there are no functorial projections to cut metrics, or even to tree metrics. Finally, although we focus less on the construction of clustering methods (clustering domains) derived from injective envelopes, we lay out some preliminary results, that hopefully will give a feel for how the third leg of the stool comes into play.

1. INTRODUCTION

Problems surrounding data clustering have been studied extensively over the last forty years. Clustering stands as an important tool for analyzing and revealing the often hidden structure in data (and in today's big data) coming from fields as diverse as biology, psychology, machine learning, sociology, image understanding, and chemistry. Among the earliest systematic treatment of clustering theory was that of Jardine and Sibson in 1971 [30]. They laid out important desiderata for overlapping clustering methods and provided a relatively efficient algorithm for their so-called B_k clustering which allowed overlapping clusters with no more than $k - 1$ points in any overlap. Since then, there have been several distinct directions of research in clustering theory, with only modest linkage between the methods of researchers pursuing different paths.

The classical work of Jardine and Sibson was followed by other similarly comprehensive works such as Everitt [22]. Further theoretical work on these mostly classical methods was also done by Kleinberg [31] and Carlsson and Mémoli [8]. Kleinberg in particular showed the incompatibility of a relatively simple set of desirable axioms for any partition based clustering method. Carlsson and Mémoli in turn introduced categorical language into partition based clustering and showed that single-linkage was (up to a scaling) the only functorial method satisfying all

their axioms (which included the notions of representability and excisiveness) in the category of finite metric spaces with non-expansive maps.

In another direction, work on computing phylogenetic trees inspired a seminal paper by Bandelt and Dress [3] on split decompositions of metrics. This line of research was continued with investigations into split systems and cut points of injective envelopes of metric spaces. Representative papers include [18] and [17]. While not explicitly clustering methods, these methods are quite similar in spirit to stratified/hierarchical clustering schemes. In this genre, we might also add the classification of the injective envelopes of six-point metric spaces by Sturmfels and Yu [36]. Bandelt and Dress have also had a large influence on another field as a result of their work on weak hierarchies [2, 4]. This led to work by Diatta, Bertrand, Barthélemy, Brucker, and others on indexed set systems (see, e.g., [5], [14], [6]). Another interesting development in this area is the work by Janowitz on ordinal clustering [29].

Recently, with the emergence of the new field of topological data analysis (TDA), work has been done on topologically-based clustering methods. This includes the Mapper algorithm by Singh, Mémoli, and Carlsson [35], as well as work on persistence-based methods [10] and Reeb graphs [25].

Meanwhile, most users of clustering methods default either to a classical linkage-based clustering method (such as single-linkage or complete linkage) or to more geometrically-based methods like k -means. Unfortunately, the wide array of clustering theories has had little impact on the actual practice of clustering. Simply put, the gap between theory and efficient practice has been hard to bridge.

In this paper we draw inspiration from three of the sources mentioned above, and strive to intertwine those three threads into a consistent principled functorial theory of dissimilarity-based clustering. Those three are the overlapping clustering of Jardine and Sibson (incorporating their desiderata), the functorial approach of Carlsson and Mémoli, and the Dress school approach to clustering, via injective envelopes, which were independently discovered by Isbell and Dress. This paper intends to fuse these approaches. Our starting point is the paper of Carlsson and Mémoli [9] which introduces the idea of viewing clustering methods as functors from a category of metric spaces to a category of clusters. Many desirable properties of a clustering method are subsumed in functoriality when the morphisms are properly chosen. Here the relevant morphisms under which the particular method is functorial can be viewed as giving restrictions on the allowable data processing operations—restrictions that impose consistency constraints across related data sets. One of our first goals is to extend their theory of functorial clustering schemes to methods that allow overlapping clusters in the spirit of Jardine and Sibson, and in so doing obviate some of the unpleasant effects of chaining occurring in some linkage-based methods. Rather than relying on chaining to overcome certain technical problems, we accept overlapping clusters. This leads to a much richer set of possible clustering algorithms.

Finally, although in this paper we focus less on the construction of clustering methods (clustering domains) derived from injective envelopes, we do in the final section lay out some preliminaries, that hopefully will enable the reader to get a feel for how the geometry of injective envelopes comes into play, and entice that reader into looking at our forthcoming paper [12] where these ideas are fleshed out. In addition, lest the reader think we are all theory and no practice, we mention our

ongoing algorithmic work on efficient implementations of some of these clustering schemes, along the lines of what has already been done for q -metrics by Segarra *et al.* [33] and for dithered maximal linkage clustering by Gama *et al.* [23].

1.1. Weight Categories.

Definition 1. Let **Weight** be the category of **finite sets with weights**, whose objects have the form (X, u) with X a finite non-empty set and u a symmetric non-negative map $u: X \times X \rightarrow \mathbb{R}$, $(x, y) \mapsto u_{xy}$ satisfying $u_{xx} = 0$ for all $x \in X$. A morphism $f: (X, u) \rightarrow (Y, v)$ is a set map $f: X \rightarrow Y$ such that $v_{f(x)f(x')} \leq u_{xx'}$; these will be referred to as *non-expansive maps*.

For a fixed set X , we can define a local order structure by pointwise dominance on the full subcategory **Weight** $_X$ of **Weight** consisting of weights on X . In order to simplify notation, when the underlying set X is fixed we will often refer to the object $(X, u) \in \mathbf{Weight}_X$ only by the weight function u and state that $u \in \mathbf{Weight}_X$. The set of objects of **Weight** $_X$ is a partially ordered set (poset) with the ordering given by

$$u \leq v \text{ if } u_{xy} \leq v_{xy} \text{ for all } x, y \in X.$$

It will be convenient to denote, for any subset U of the objects of **Weight** $_X$,

$$U \downarrow := \left\{ w \in \mathbf{Weight}_X \mid \exists u \in U \ w \leq u \right\}$$

and in the case of the singleton set $\{u\}$, we will often write $u \downarrow$ for $\{u\} \downarrow$. For any subcategory **C** of **Weight**, we will use the analogous notation **C** $_X$ for **C** \cap **Weight** $_X$, which on objects will be the intersection of the objects of **C** and **Weight** $_X$ with the morphisms of **C**. Also, for any set map $f: X \rightarrow Y$ and $w \in \mathbf{Weight}_Y$ we define $f^*(w) \in \mathbf{Weight}_X$ to be the pullback of the weight w to X , explicitly, $f^*(w)_{xy} := w_{f(x)f(y)}$. This notation allows another perspective on morphisms in **Weight**: the set map $f: X \rightarrow Y$ induces a morphism $(X, u) \rightarrow (Y, v)$ in **Weight** if and only if $f^*(v) \leq u$.

Definition 2. By a **weight category** we mean any full subcategory **C** of **Weight** such that $f^*\mathbf{C}_Y \subseteq \mathbf{C}_X$ for any set map $f: X \rightarrow Y$. We say that **C** is closed under pullbacks, and will sometimes refer to this as the pullback property. Note that this condition is only on the objects of the category **C**.

Note that the category **Met** of metric spaces and non-expansive maps, first considered by Isbell [28], is a weight category in this sense. In this exposition, by metric we will always mean semimetric; *i.e.*, we will not require that distinct points have nonzero distance.

Remark 3. Since the pullback property allows arbitrary set maps, any permutation mapping $f: X \rightarrow X$ satisfies $f^*\mathbf{C}_X = \mathbf{C}_X$. In other words, weight categories are invariant under permutations.

Definition 4. A functor $\mathcal{G}: \mathbf{C} \rightarrow \mathbf{D}$ between weight categories will be said to be **fibered**, if the diagram

$$\begin{array}{ccc}
\mathbf{C} & \xrightarrow{\mathcal{G}} & \mathbf{D} \\
& \searrow \mathcal{F}_{\mathbf{C}} & \downarrow \mathcal{F}_{\mathbf{D}} \\
& & \mathbf{Set}
\end{array}$$

commutes, where $\mathcal{F}_{\mathbf{C}}$ and $\mathcal{F}_{\mathbf{D}}$ are the respective forgetful functors. For such a functor \mathcal{G} and any finite set X , we have an associated functor $\mathcal{G}_X: \mathbf{C}_X \rightarrow \mathbf{D}_X$. In other words, on objects, fibered functors fix the underlying set and on morphisms, fix the underlying set map. Intuitively, a fibered functor acts on weighted spaces only by deforming the weight structure.

Remark 5. This entire exposition could be reformulated in the setting of fibered categories and functors, where the fiber over a fixed set has the structure of a directed complete poset, but we did not view the additional advantage in abstraction and simplicity of definitions worth the overhead cost of framing the work in those terms. The fibration of a weight category over **Set** is straightforward. In Section 4.3, we give a hint at where this generality would be important: we explore a subcategory of **Weight** with a restricted set of morphisms. The important thing in this type of setting is that the cartesian morphisms in each fiber provide this directed complete poset structure. For **Weight**, the categories considered in Section 4.3, as well as for the categories considered by Carlsson and Mémoli [9], this amounts to noting that the morphisms with underlying map the identity are cartesian.

2. CLUSTERING WITH OVERLAPS: SIEVES

Let **Cov** be the category of coverings, with an object a pair (X, C) consisting of a set X together with a cover C of X . A morphism $f: (X, C) \rightarrow (Y, D)$ is a set map $f: X \rightarrow Y$ such that the cover $f^{-1}(D)$ of X is refined by C .

Definition 6. Given a non-empty finite set X we define a **non-nested flag cover** of X to be a cover C of X additionally satisfying:

- (i) for all $A, B \in C$ with $A \subseteq B$, we have $A = B$;
- (ii) the abstract simplicial complex with vertices corresponding to elements of X and faces all subsets of the sets in C is a flag complex.

The category $\mathbf{Cov}_{\mathbf{P}}$ is then the full subcategory of **Cov** where the covers are required to be non-nested flag covers.

Remark 7. For a non-nested flag cover C of X , observe that the sets in C provide the maximal simplices in the simplicial complex defined in (ii) above. Given a fixed set X , we will denote the full subcategory of $\mathbf{Cov}_{\mathbf{P}}$ consisting of covers on X by $\mathbf{Cov}_{\mathbf{P}}(X)$. Notice that, in particular, every partition of X is a flag cover of X .

Remark 8. The inclusion functor $\mathbf{Cov}_{\mathbf{P}} \rightarrow \mathbf{Cov}$ has a left adjoint given by flagification. Any cover \tilde{C} of X can be flagified to a non-nested flag cover C in a minimal way, where \tilde{C} will refine C and C will refine any other non-nested flag cover which \tilde{C} refines. Algorithmically, this is very simple and amounts to adding in any subsets required by the flag condition and then removing the nested sets from the cover. In general, for a set map $f: X \rightarrow Y$ and a flag cover D of Y , the cover $f^{-1}(D)$

of X might contain nesting. The condition given above for morphisms is correct, however, in the sense that if C is a flag cover on X , then $f^{-1}(D)$ is refined by C if and only if the flagification of $f^{-1}(D)$ is refined by C .

Definition 9. A **sieve** on X is a function $\theta: \mathbb{R}_{\geq 0} \rightarrow \mathbf{Cov}_{\mathbf{p}}(X)$ such that

- (i) If $t_1 < t_2$ then $\theta(t_1)$ refines $\theta(t_2)$
- (ii) For any t , there is an $\varepsilon > 0$ such that $\theta(t') = \theta_X(t)$ for all $t' \in [t, t + \varepsilon)$
- (iii) There exists $t \in \mathbb{R}_{\geq 0}$ such that $\theta(t)$ is the trivial cover $\{X\}$.

We say that a sieve θ on S is **proper**, if $\theta(0)$ is the covering by singletons. In order to curb the proliferation of parentheses, we will also use the notation θ_t for $\theta(t)$.

Sieves are the obvious generalization of *dendograms*, which satisfy the same conditions but take values in the set of partitions of X rather than coverings.

Definition 10. We define the **category of sieves** **Sieve** as the category of pairs (X, θ) , where $\theta: \mathbb{R}_{\geq 0} \rightarrow \mathbf{Cov}_{\mathbf{p}}(X)$ is a sieve on X . The morphisms in **Sieve** are an extension of the morphisms of $\mathbf{Cov}_{\mathbf{p}}$; that is, a set map $f: X \rightarrow Y$ is a morphism of sieves $(X, \theta) \rightarrow (Y, \psi)$ if for every $t \in \mathbb{R}_{\geq 0}$, θ_t refines $f^{-1}(\psi_t)$. In other words, for each $t \in \mathbb{R}_{\geq 0}$, we have a functor $t: \mathbf{Sieve} \rightarrow \mathbf{Cov}_{\mathbf{p}}$. Note that just as in Remark 8, $f^{-1}(\psi_t)$ might contain nesting, but nonetheless, this condition is correct.

Just as in **Weight**, the categories $\mathbf{Cov}_{\mathbf{p}}$ and **Sieve** are also fibered over **Set** in a natural way. We will restrict to fibered functors in our discussion, which in particular ensures that underlying sets and mappings are fixed by the functors. Any fibered functor with values in **Sieve** then provides a hierarchical method of clustering that respects the constraints imposed by the morphisms in **Weight** and produces potentially overlapping clusters at any fixed scale. For clarity in the exposition, we make the following definition.

Definition 11. A **sieving functor** is a fibered functor $\mathbf{Weight} \rightarrow \mathbf{Sieve}$. One of the primary purposes of this paper is to characterize a particularly “nice” class of sieving functors, in a way that will be made precise below (see Theorem 28).

Example 12 (Rips Sieving). Let $(X, u) \in \mathbf{Weight}$ and let $\delta \geq 0$. Recall that the Rips complex at resolution δ on (X, u) is the abstract simplicial complex $K_u(\delta)$ with vertex set X , where a subset $A \subseteq S$ forms a face of $K_u(\delta)$ if and only if $\text{diam}(A) \leq \delta$. If $M_u(\delta)$ is the collection of maximal simplices of $K_u(\delta)$, then $M_u(\delta)$ forms a non-nested flag cover of X . Define $\mathcal{R}: \mathbf{Weight} \rightarrow \mathbf{Sieve}$ by

$$\mathcal{R}(X, u)_t = (X, M_u(t)),$$

which we will call the Rips sieving functor (elsewhere, including in [11], this functor is denoted \mathcal{ML} and called maximal-linkage). On morphisms, \mathcal{R} sends a non-expansive map of weight spaces to the same underlying set function. It is straightforward to see that this gives a **Sieve**-morphism.

Proposition 13. *The category **Weight** is equivalent to **Sieve**.*

Proof. Given a sieve (X, θ) , define a weight space (X, u_θ) pointwise with $x, y \in X$ by

$$u_\theta(x, y) = \min \left\{ t \in \mathbb{R}_{\geq 0} \mid \exists A \in \theta_t, \{x, y\} \subseteq A \right\}.$$

At first glance, it might appear that we need to consider the infimum rather than the minimum, but condition (ii) in Definition 9 ensures that u_θ is well-defined. This

assignment extends to a functor $\mathcal{J}: \mathbf{Sieve} \rightarrow \mathbf{Weight}$, and it is straightforward to check that $\mathcal{J} \circ \mathcal{R}, \mathcal{R} \circ \mathcal{J}$ are the respective identity functors. \square

Remark 14. This proposition shows that Rips sieving is the canonical way of producing sieves from weight spaces, in the sense that a weight space can be recovered from its Rips sieve. The result extends the connection noted by Carlsson–Mémoli between the category of ultrametrics (Definition 33) and the category of dendrograms; indeed our proof is a reworking of theirs in this new context (see [8], Theorem 9). From this perspective, a natural question is how to characterize sieving functors that factor through Rips sieving. Theorem 28 below shows one result in this direction.

Example 15 (Single-Linkage Sieving). Let $(X, u) \in \mathbf{Weight}$ and $K_u(\delta)$ as in Example 12. In this case, rather than look at maximal simplices, let $C_u(\delta)$ be the connected components of $K_u(\delta)$. Then we can define the single-linkage sieving functor $\mathcal{SL}: \mathbf{Weight} \rightarrow \mathbf{Sieve}$ by

$$\mathcal{SL}(X, u)_t = (X, C_u(t)),$$

with morphisms being sent to the same underlying set maps. This is equivalent to the many alternative well-known characterizations of single-linkage clustering, for instance in [34, 24, 9].

Example 16 (Čech Sieving). If $(X, u) \in \mathbf{Weight}$, define a graph $G_u(\delta)$ with vertices the elements of X and an edge xy if there exists an element $z \in X$ with $u(x, z), u(z, y) \leq \delta$. We then define $\check{\mathcal{C}}: \mathbf{Weight} \rightarrow \mathbf{Sieve}$ to be the assignment which sends any $t \in \mathbb{R}_{\geq 0}$ to the set of maximal cliques of $G_u(t)$. Again in this case, morphisms are sent to the same underlying set maps.

Many other examples can be derived from the clustering functors introduced in our work with Hansen in [11].

3. CLUSTERING DOMAINS

Definition 17. We say that a weight category \mathbf{D} is a **clustering domain** if for any finite non-empty set X and any $w \in \mathbf{Weight}_X$, the set $\mathbf{D}_X \cap w \downarrow$ is non-empty and sup-closed in \mathbf{Weight}_X . That is,

$$(X, \sup\{u \in \mathbf{D}_X \cap w \downarrow\}) \in \mathbf{D},$$

for all $w \in \mathbf{Weight}_X$.

It will often be convenient to satisfy a stronger (but more easily verified) condition related to the topology on \mathbf{Weight} induced by the identification of the objects of \mathbf{Weight}_X with the nonnegative real orthant $\mathbb{R}_{\geq 0}^{\binom{|X|}{2}}$.

Proposition 18. *Let \mathbf{C} be a weight category. If the set of objects of \mathbf{C}_X is closed as a subset of $\mathbb{R}_{\geq 0}^{\binom{|X|}{2}}$ and \mathbf{C}_X is closed under taking finite maxima, then \mathbf{C} is a clustering domain.*

Proof. If $w \in \mathbf{Weight}_X$ and $S \subset \mathbf{C}_X$ satisfies $d \leq w$ for all $d \in S$, then we see that $v := \sup_{d \in S} (d)$ satisfies $v \leq w$ and can consider the following procedure. For every pair $x, y \in X$ and $\varepsilon > 0$ find $u_{x,y}^\varepsilon \in S$ such that $v(x, y) - u_{x,y}^\varepsilon(x, y) < \varepsilon$; then $u^\varepsilon := \max_{x,y \in X} u_{x,y}^\varepsilon \in \mathbf{C}_X$ and satisfies $\|v - u^\varepsilon\|_\infty < \varepsilon$, and so we conclude that v is in the closure of \mathbf{C}_X . \square

3.1. Projections. The reason to define clustering domains is that, in some sense, clustering maps “live” on them.

Definition 19 (Canonical Projection). Given a clustering domain \mathbf{D} , the **canonical projection** $\mathcal{P}_{\mathbf{D}}: \mathbf{Weight} \rightarrow \mathbf{Weight}$ is the fibered functor defined by

$$\mathcal{P}_{\mathbf{D}}(X, w) = (X, \sup\{u \in \mathbf{D}_X \cap w \downarrow\}).$$

Recall that the fibered condition implies that $\mathcal{P}_{\mathbf{D}}$ sends a morphism in \mathbf{Weight} to the morphism with the same underlying set map. Also $(X, \sup\{u \in \mathbf{D}_X \cap w \downarrow\})$ is an object in \mathbf{D} since \mathbf{D} is a clustering domain.

We need to verify the assertions made in this definition regarding properties of $\mathcal{P}_{\mathbf{D}}$.

Proposition 20 (Properties of the Canonical Projection). *Suppose \mathbf{D} is a clustering domain and $\mathcal{P} = \mathcal{P}_{\mathbf{D}}$ is the associated canonical projection. Then \mathcal{P} is a fibered endofunctor of \mathbf{Weight} satisfying the additional properties:*

- (1) $\mathcal{P} \circ \mathcal{P} = \mathcal{P}$,
- (2) $\mathcal{P}_X w \leq w$ for all sets X and $w \in \mathbf{Weight}_X$.

Proof. Since the identities

$$\mathcal{P}(\text{id}_{(X,w)}) = \text{id}_{\mathcal{P}(X,w)}, \quad \mathcal{P}(g \circ f) = \mathcal{P}(g) \circ \mathcal{P}(f)$$

are immediate from the definition, one is left only to verify that $\mathcal{P}(f)$ is a non-expansive map whenever $f: (X, u) \rightarrow (Y, v)$ is. Equivalently, we need to show that $f^*(\mathcal{P}v) \leq \mathcal{P}u$ holds whenever $f^*(v) \leq u$. By definition, $\mathcal{P}v \leq v$ and we have:

$$\mathcal{P}v \leq v \Leftrightarrow f^*(\mathcal{P}v) \leq f^*(v) \Rightarrow f^*(\mathcal{P}v) \leq u \Leftrightarrow f^*(\mathcal{P}v) \in u \downarrow$$

Thus, $f^*(\mathcal{P}v) \in u \downarrow \cap f^*(\mathbf{D}_Y) \subseteq u \downarrow \cap \mathbf{D}_X$, by the definition of a clustering domain. Finally, $f^*(\mathcal{P}v) \leq \mathcal{P}u$ by the definition of $\mathcal{P}u$. \square

In an attempt to construct more general sieving functors one might seek to introduce the following definition:

Definition 21 (Projection). Let \mathbf{C} be a weight category. We say that a fibered functor $\mathcal{P}: \mathbf{C} \rightarrow \mathbf{C}$ is a **projection**, if it satisfies:

- **Idempotency** $\mathcal{P} \circ \mathcal{P} = \mathcal{P}$;
- **Contraction** $\mathcal{P}_X w \leq w$ for all sets X and $w \in \mathbf{C}_X$.

Remark 22. The fact that \mathcal{P} is a fibered functor implies that for a fixed set X , the self-maps $\mathcal{P}_X: \mathbf{Weight}_X \rightarrow \mathbf{Weight}_X$ are order-preserving. Indeed, for $w_1, w_2 \in \mathbf{Weight}_X$ one has:

$$\begin{aligned} w_1 \leq w_2 &\Leftrightarrow \text{id}_X \in \text{Hom}_{\mathbf{Weight}}((X, w_2), (X, w_1)) \\ &\Rightarrow \text{id}_X = \mathcal{P}(\text{id}_X) \in \text{Hom}_{\mathbf{Weight}}((X, \mathcal{P}_X w_2), (X, \mathcal{P}_X w_1)) \\ &\Leftrightarrow \mathcal{P}_X w_1 \leq \mathcal{P}_X w_2 \end{aligned}$$

The following result was initially observed for single linkage hierarchical clustering in [38], however its categorification and generalization are new:

Proposition 23 (Uniqueness of Projections). *Let $\mathbf{C} \subseteq \mathbf{Weight}$ be a weight category. Then every weight category $\mathbf{D} \subseteq \mathbf{C}$ admits at most one projection $\mathcal{P}: \mathbf{C} \rightarrow \mathbf{C}$ whose image coincides with \mathbf{D} .*

Proof. Suppose \mathcal{P}, \mathcal{Q} are arbitrary projections with image \mathbf{D} . Fixing a non-empty finite set X and applying the idempotency requirement we observe that, for all $w \in \mathbf{C}_X$:

$$w \in \mathbf{D}_X \Leftrightarrow \mathcal{P}_X w = w \Leftrightarrow \mathcal{Q}_X w = w$$

We conclude that \mathcal{P}_X fixes $\mathcal{Q}_X w$ and \mathcal{Q}_X fixes $\mathcal{P}_X w$ for all $w \in \mathbf{C}_X$. From Remark 22 we then obtain, for all $w \in \mathbf{C}_X$:

$$\mathcal{P}_X w \leq w \Rightarrow \mathcal{P}_X w = \mathcal{Q}_X \mathcal{P}_X w \leq \mathcal{Q}_X w$$

We conclude that $\mathcal{P}_X w \leq \mathcal{Q}_X w$ for all $w \in \mathbf{C}_X$. By a symmetric argument, the reverse inequality holds true as well and we have shown that \mathcal{P} and \mathcal{Q} coincide on \mathbf{C} . \square

Corollary 24. *Suppose \mathbf{C} is a weight category containing a clustering domain \mathbf{D} . Then \mathbf{C} admits one and only one projection with image \mathbf{D} : the restriction of the canonical projection $\mathcal{P}_{\mathbf{D}}$ to \mathbf{C} .* \square

This last corollary emphasizes that the existence of a clustering projection must be characterized in terms of its image category. We have:

Theorem 25 (Existence of Projections). *Let \mathbf{D} be full subcategory of \mathbf{Weight} . Then \mathbf{D} is a clustering domain if and only if it is the image of a projection $\mathcal{P}: \mathbf{Weight} \rightarrow \mathbf{Weight}$.*

Proof. If \mathbf{D} is a clustering domain then it is the image of the canonical projection $\mathcal{P} = \mathcal{P}_{\mathbf{D}}$. Conversely, assume \mathbf{D} is the image of a clustering projection $\mathcal{P}: \mathbf{Weight} \rightarrow \mathbf{Weight}$ and let us prove it is a clustering domain.

First, we need to see that \mathbf{D} is a weight category. Suppose $f: X \rightarrow Y$ is a set map. For $w \in \mathbf{D}_Y$, consider $v = f^*w \in \mathbf{Weight}_X$: then $f: (X, v) \rightarrow (Y, w)$ is a morphism and hence $f: (X, \mathcal{P}_X v) \rightarrow (Y, \mathcal{P}_Y w) = (Y, w)$ is a morphism as well; equivalently, $v = f^*w \leq \mathcal{P}_X v$. Thus, by the contraction property of \mathcal{P} we have $v = \mathcal{P}_X v \in \mathbf{D}_X$, as desired.

Note that for any set X and $w \in \mathbf{Weight}_X$, the set $\mathbf{D}_X \cap w \downarrow$ is nonempty since in particular, this contains $\mathcal{P}w$. It remains to prove that \mathbf{D}_X is sup-closed. Let U be the set of objects in $\mathbf{D}_X \cap w \downarrow$. If $v = \sup\{u \in U\}$ then v is the minimal weight in \mathbf{Weight}_X satisfying $v \geq u$ for all $u \in U$. By functoriality, we see that $\mathcal{P}v \geq \mathcal{P}u = u$, so we must have $\mathcal{P}v \geq v$. Since \mathcal{P} is a projection, and in particular, satisfies the contraction property, $\mathcal{P}v \leq v$ as well and so $v \in \mathbf{D}_X$, as required. \square

Remark 26. Let $\mathbf{C} \subset \mathbf{D}$ be clustering domains. Proposition 23 implies that we have a unique factorization:

$$\begin{array}{ccc} & \mathbf{Weight} & \\ \mathcal{P}_{\mathbf{D}} \downarrow & \searrow \mathcal{P}_{\mathbf{C}} & \\ \mathbf{D} & \xrightarrow{\mathcal{P}_{\mathbf{C}}|_{\mathbf{D}}} & \mathbf{C} \end{array}$$

where $\mathcal{P}_{\mathbf{C}}$ and $\mathcal{P}_{\mathbf{D}}$ are the respective projections associated with \mathbf{C} and \mathbf{D} . Note that the essential portions of the proof of the proposition is based on the order structure of \mathbf{Weight}_X , and consequently, this factorization holds even when replacing \mathbf{Set}

with a category with fewer morphisms (see Remark 5), such as only injective maps or only surjective maps.

Definition 27. Let $\mathcal{C}: \mathbf{Weight} \rightarrow \mathbf{Sieve}$ be a sieving functor and recall the functor $\mathcal{J}: \mathbf{Sieve} \rightarrow \mathbf{Weight}$ from the proof of Proposition 13. Then \mathcal{C} will be called **stationary** if $\mathcal{J} \circ \mathcal{C}$ is a projection.

We are now able to state a theorem characterizing stationary sieving functors.

Theorem 28. *There is a bijective correspondence between the collection of stationary sieving functors and the collection of clustering domains: every stationary sieving functor factors uniquely through a clustering projection and the Rips sieving functor \mathcal{R} restricted to a clustering domain. Pictorially, there is a commutative diagram*

$$\begin{array}{ccc} \mathbf{Weight} & & \\ \mathcal{P}_{\mathbf{D}_c} \downarrow & \searrow \mathcal{C} & \\ \mathbf{D}_c & \xrightarrow{\mathcal{R}|_{\mathbf{D}_c}} & \mathbf{Sieve}. \end{array}$$

Proof. If $\mathcal{C}: \mathbf{Weight} \rightarrow \mathbf{Sieve}$ is a stationary sieving functor, then $\mathcal{P} = \mathcal{J} \circ \mathcal{C}$ is a projection and so from Theorem 25 we see that the image of \mathcal{P} is a clustering domain. The theorem then follows from Proposition 13. \square

In our first look at sieving functors, we defined the Čech sieve of Example 16, which is not a stationary sieving functor: after sufficiently many applications it gives the result of the single-linkage sieve. Indeed this is generally the case. In our previous paper with Hansen [11], we exhibited many different sieving functors that fail to be stationary. For a fixed weight space, however, iterated applications of any sieving functor (interleaving applications of the functor \mathcal{J}) will eventually lead to the same result as applying a stationary sieving functor. An interesting generalization of the work here would be to explore more complex characterizations that would include non-stationary sieving functors.

3.2. Single-linkage as a projection. In this section, we illustrate the relationship between projections and clustering domains by revisiting the well-known single-linkage clustering method from a new perspective. More details than necessary are given in order to provide a template for one way of using the results in the previous section to construct a clustering projection.

Let $(X, u) \in \mathbf{Weight}$ and $x_1, x_2 \in X$ with $x_1 \neq x_2$. Also, let Δ_2^ε be the two point space $\{1, 2\}$ with distance ε between the two points. Then define $\mathcal{E}_u(x_1, x_2)$ to be the largest ε for which there exists a morphism (non-expansive map) $f: (X, u) \rightarrow \Delta_2^\varepsilon$, with $f(x_i) = i$ for $i = 1, 2$. For $x_1 = x_2$, define $\mathcal{E}_u(x_1, x_2) = 0$. Note that if the smallest distance between any two points in X is u_{sep} (the so-called separation of X), then $\mathcal{E}_u(x_1, x_2) \geq u_{\text{sep}}$ since we can map X to $\Delta_2^{u_{\text{sep}}}$ arbitrarily and have a non-expansive map. Moreover, notice that $\mathcal{E}_u(x_1, x_2) \leq u(x_1, x_2)$ since f is required to be non-expansive. This means that the identity set map on X gives a **Weight**-morphism $(X, u) \rightarrow (X, \mathcal{E}_u)$, or in the language of Definition 21, \mathcal{E} is a contraction.

Definition 29. Recall that a **cut metric** (see [13] for a detailed treatment) on any non-empty set Y is a non-negatively weighted finite combination of cuts. A cut is a metric given by a subset $A \subset X$ and its complement A^c , with the associated metric δ_A on X is defined as

$$\delta_A(y_1, y_2) = \begin{cases} 0 & \text{if } y_1, y_2 \text{ are both in } A \text{ or both in } A^c \\ 1 & \text{otherwise.} \end{cases}$$

Notice that $\delta_A = \delta_{A^c}$ and the trivial cut $\{X, \emptyset\}$ produces the zero metric.

Now for any fixed $x_1, x_2 \in X$, and any morphism $f: (X, u) \rightarrow \Delta_2^{\mathcal{E}_u(x_1, x_2)}$ with $f(x_i) = i$ for $i = 1, 2$, the morphism f defines a bi-partition of X into $X_1 = f^{-1}(1)$ and $X_2 = f^{-1}(2)$ with $x_1 \in X_1$ and $x_2 \in X_2$. The morphism f will then factor through X with the cut metric $\delta_{\{X_1\}} = \delta_{\{X_2\}}$ scaled by $\mathcal{E}_u(x_1, x_2)$:

$$\begin{array}{ccc} (X, u) & \xrightarrow{f} & \Delta_2^{\mathcal{E}_u(x_1, x_2)} \\ & \searrow & \uparrow \\ & & (X, \mathcal{E}_u(x_1, x_2)\delta_{\{X_1\}}). \end{array}$$

Notice that $\mathcal{E}_u(x_1, x_2)\delta_{\{X_1\}}$ is just the pullback of the metric $\Delta_2^{\mathcal{E}_u(x_1, x_2)}$ by f . In general, we might have several morphisms $f: (X, u) \rightarrow \Delta_2^{\mathcal{E}_u(x_1, x_2)}$ which give different bi-partitions.

We now show that the mapping \mathcal{E} taking (X, u) to (X, \mathcal{E}_u) is a projection as described in Definition 21.

Proposition 30. *The mapping $\mathcal{E}: \mathbf{Weight} \rightarrow \mathbf{Weight}$ is a fibered functor with the contraction property, where \mathcal{E} takes a morphism $f: (X, u) \rightarrow (Y, v)$ in \mathbf{Weight} to the same underlying set map.*

Proof. Since \mathcal{E} takes morphisms to the same underlying set map, it is fibered, and as has already been observed above, \mathcal{E} is a contraction. Hence the only thing requiring proof is that the map $\mathcal{E}(g): (X, \mathcal{E}_u) \rightarrow (Y, \mathcal{E}_v)$ is a **Weight**-morphism. For this, we simply need to see that for every pair of points $x_1, x_2 \in X$ we have $\mathcal{E}_u(x_1, x_2) \geq \mathcal{E}_v(g(x_1), g(x_2))$. But we have morphisms g and f :

$$(X, u) \xrightarrow{g} (Y, v) \xrightarrow{f} \Delta_2^{\mathcal{E}_v(f(x_1), f(x_2))}$$

whose composition $f \circ g$ is a morphism sending x_1 and x_2 to different points in $\Delta_2^{\mathcal{E}_v(f(x_1), f(x_2))}$. Since $\mathcal{E}_u(x_1, x_2)$ is defined to be the maximum ε satisfying this property, we must have $\mathcal{E}_u(x_1, x_2) \geq \mathcal{E}_v(g(x_1), g(x_2))$. \square

Proposition 31. *The functor \mathcal{E} is idempotent, i.e. $\mathcal{E} \circ \mathcal{E} = \mathcal{E}$.*

Proof. Since \mathcal{E} sends morphisms to the same underlying set map, we only need to check that \mathcal{E} is idempotent on the objects of **Weight**. Notice that for any $\varepsilon > 0$, we have $\mathcal{E}(\Delta_2^\varepsilon) = \Delta_2^\varepsilon$. Thus for any morphism $f: (X, u) \rightarrow \Delta_2^\varepsilon$ separating a pair of points $x_1, x_2 \in X$, we have a corresponding morphism $\mathcal{E}(f): (X, \mathcal{E}_u) \rightarrow \Delta_2^\varepsilon$ also separating x_1, x_2 . Combining this with the fact that $\mathcal{E}^2(x_1, x_2) \leq \mathcal{E}(x_1, x_2)$, we get equality. \square

The previous two propositions then show that \mathcal{E} is a projection and so by Proposition 23, the functor \mathcal{E} is the canonical projection for its image, which must be a

clustering domain. The remaining question is how to characterize this clustering domain. The following lemma will be useful toward this end.

Lemma 32. *Let \mathbf{D} be a clustering domain containing Δ_2^ε for all $\varepsilon \geq 0$. Then $\text{im}(\mathcal{E}) \subseteq \mathbf{D}$.*

Proof. Let $\mathcal{P}_{\mathbf{D}}: \mathbf{Weight} \rightarrow \mathbf{Weight}$ be the canonical projection with image \mathbf{D} and $(X, u) \in \mathbf{Weight}$. Then for any $x_1, x_2 \in X$, we have

$$\mathcal{E}_u(x_1, x_2) \leq \mathcal{P}_{\mathbf{D}}(u)(x_1, x_2) \leq u(x_1, x_2)$$

since $f: (X, u) \rightarrow \Delta_2^{\mathcal{E}_u(x_1, x_2)}$ yields a morphism $\mathcal{P}_{\mathbf{D}}(f): (X, \mathcal{P}_{\mathbf{D}}(u)) \rightarrow \Delta_2^{\mathcal{E}_u(x_1, x_2)}$. Proposition 31 then implies that $\mathcal{E} \circ \mathcal{P}_{\mathbf{D}} = \mathcal{E}$. Replacing u with \mathcal{E}_u in the above reasoning, we also get

$$\mathcal{E}_u(x_1, x_2) \leq \mathcal{P}_{\mathbf{D}}(\mathcal{E}_u)(x_1, x_2) \leq \mathcal{E}_u(x_1, x_2),$$

and we see that $\mathcal{P}_{\mathbf{D}}$ fixes $\text{im}(\mathcal{E})$. \square

For further discussion of this point in a different context, see [11], Theorem 8. Here Lemma 32 extends the concept of the referenced theorem to the hierarchical setting, where the infrastructure of clustering domains and projections affords a succinct statement. In order to state the next result, we need to recall the definition of an ultrametric.

Definition 33. An **ultrametric** on X is a metric (X, u) satisfying a stronger version of the triangle inequality: for any x, y, z in X , we have

$$u(x, z) \leq \max\{u(x, y), u(y, z)\}.$$

The category of ultrametrics **Ult** is the full subcategory of **Weight** with objects the ultrametric spaces.

Corollary 34. *For any $(X, u) \in \mathbf{Weight}$, the weight (X, \mathcal{E}_u) is an ultrametric.*

Proof. Using Lemma 32, we need only to note that the category of ultrametrics contains all two point weight spaces. Alternatively, one could verify the ultrametric inequality directly as follows. Let $x_1, x_2 \in X$ with $x_1 \neq x_2$, let $y \in X$ and let $f: (X, u) \rightarrow \Delta_2^{\mathcal{E}_u(x_1, x_2)}$ be a **Weight**-morphism that separates x_1 and x_2 , so that $f(x_i) = i$. Now, if $y \in f^{-1}(1)$ then $\mathcal{E}_u(x_2, y) \geq \mathcal{E}_u(x_1, x_2)$ by construction. Likewise, if $y \in f^{-1}(2)$ then $\mathcal{E}_u(x_1, y) \geq \mathcal{E}_u(x_1, x_2)$. Since y is in one or the other, we must have $\mathcal{E}_u(x_1, x_2) \leq \max\{\mathcal{E}_u(x_1, y), \mathcal{E}_u(x_2, y)\}$, as desired. \square

Proposition 35. *The image of \mathcal{E} is precisely the category **Ult** of ultrametrics.*

Proof. Following the previous corollary, we need to see that every ultrametric (X, u) is in the image of \mathcal{E} (equivalently, that u is fixed by \mathcal{E}). Given an ultrametric (X, u) , and $x_1, x_2 \in X$ with $x_1 \neq x_2$, it suffices to exhibit a non-expansive map from (X, u) to $\Delta_2^{u(x_1, x_2)}$. If $u(x_1, x_2) = 0$ any map will do, provided x_i is sent to i .

The ultrametric property allows us to define an equivalence relation on X by $y_1 \sim y_2$ if $u(y_1, y_2) < u(x_1, x_2)$. We can then construct a **Weight**-morphism $f: (X, u) \rightarrow \Delta_2^{u(x_1, x_2)}$ as follows. Since x_1 and x_2 are in different equivalence classes we map everything equivalent to x_1 to 1 and everything equivalent to x_2 to 2. The remaining equivalence classes can be mapped arbitrarily as long as the entire equivalence class is sent to the same point in $\Delta_2^{u(x_1, x_2)}$. Because points in different equivalence classes are at least $u(x_1, x_2)$ apart, we see that f is non-expansive. \square

Since we know that single-linkage clustering factors through the unique projection to the category of ultrametrics (see Carlsson and Mémoli [9]), we now see that, in fact, the functor \mathcal{E} is another description of this projection. In other words, given a weight (X, u) , the weight \mathcal{E}_u is the maximal ultrametric under u . Moreover, we can compose with the Rips sieving functor $\mathcal{R}: \mathbf{Weight} \rightarrow \mathbf{Sieve}$ and recover the single-linkage sieve:

$$\begin{array}{ccc}
 \mathbf{Weight} & & \\
 \downarrow \mathcal{E} & \searrow \mathcal{SL} & \\
 \mathbf{Ult} & \xrightarrow{\mathcal{R}} & \mathbf{Sieve}
 \end{array}$$

The construction described in this section gives us a characterization of the maximal ultrametric \mathcal{E}_u on X underneath a given weight u :

$$\mathcal{E}_u(x_1, x_2) = \max_{\substack{A \subset X \\ x_1 \in A \\ x_2 \in A^c}} \min_{\substack{y_1 \in A \\ y_2 \in A^c}} u(y_1, y_2),$$


where $A^c = X \setminus A$. In other words, we look at splits A, A^c of X which separate x_1 and x_2 and take one where the two sets are as far apart as possible.

Finally, the clustering domain **Ult** is characterized (by Lemma 32 and Proposition 35) as the smallest clustering domain containing all two-point spaces, allowing one to think of hierarchical single linkage clustering as *merely the coarsest* among all stationary sieving methods which agree with the Rips sieve on the two point spaces.

3.3. Examples of Clustering Domains. By enriching the setting of functorial clustering to allow non-partition based methods, we discover that there are many viable clustering domains (and hence stationary sieving functors via Theorem 28) with various potentially advantageous properties. Here we present a number of examples that help to illuminate the ideas presented in the theoretical results above and which clarify what is, or is not, a clustering domain. Once we have identified some clustering domains, additional ones can be constructed by taking intersections (by intersection, we mean the full subcategory of **Weight** with objects given by taking intersection of the objects). We can also take categories that have the relevant pull-back property under non-expansive maps, but fail to be sup closed, and “sup close” them to obtain a clustering domain. Even if the pull-back property fails, one can restrict the class of non-expansive maps to so-called admissible (non-expansive) maps where the pull-back property holds. This necessarily leads to a less restrictive notion of functoriality. We discuss all of these ideas below. However, our first example is more in the way of a counterexample.

Example 36 (Cut metrics and tree metrics). Consider the following metric d on the set $X = \{1, 2, 3, 4, 5\}$:

$$d = \begin{pmatrix} 0 & 1 & 1 & 1 & 2 \\ 1 & 0 & 2 & 2 & 1 \\ 1 & 2 & 0 & 2 & 1 \\ 1 & 2 & 2 & 0 & 1 \\ 2 & 1 & 1 & 1 & 0 \end{pmatrix}$$



where d_{ij} is the distance from i to j . This metric is the path metric on the graph pictured above where all the edges have length 1.

It is well known that (X, d) is not a cut metric (see Definition 29). Using the machinery developed in Section 3.1, we will see how this simple concrete example precludes the existence of a functorial projection to the category of cut metrics. Toward this end, consider now the metrics

$$d_1 = \begin{pmatrix} 0 & 1 & 1 & 1 & 2 \\ 1 & 0 & \frac{4}{3} & \frac{4}{3} & 1 \\ 1 & \frac{4}{3} & 0 & \frac{4}{3} & 1 \\ 1 & \frac{4}{3} & \frac{4}{3} & 0 & 1 \\ 2 & 1 & 1 & 1 & 0 \end{pmatrix} \quad \text{and} \quad d_0 = \begin{pmatrix} 0 & 1 & 1 & 1 & 0 \\ 1 & 0 & 2 & 2 & 1 \\ 1 & 2 & 0 & 2 & 1 \\ 1 & 2 & 2 & 0 & 1 \\ 0 & 1 & 1 & 1 & 0 \end{pmatrix}$$

which have cut decompositions

$$d_1 = \frac{1}{3}\delta_{\{1,2\}} + \frac{1}{3}\delta_{\{1,3\}} + \frac{1}{3}\delta_{\{1,4\}} + \frac{1}{3}\delta_{\{2,5\}} + \frac{1}{3}\delta_{\{3,5\}} + \frac{1}{3}\delta_{\{4,5\}} \\ d_0 = \delta_{\{2\}} + \delta_{\{3\}} + \delta_{\{4\}}.$$

Then clearly we have $d_0, d_1 < d$, with $d = \max\{d_0, d_1\}$.

Thus the subcategory **Cut** of **Weight** is not closed under taking max, and so is not a clustering domain. This leaves the question, however, of whether we could “sup-close” the category of cut metrics by throwing in some additional metric spaces and arrive at some proper subcategory of **Met**. The following lemma and its corollaries show that this is not the case; the reasoning used in the five-point example above is in fact general. Indeed, only finite maximums are needed to produce any metric, and the smaller category of tree metrics (for example, see [7]) is sufficient to generate all of **Met**.

Theorem 37. *Let X be a non-empty finite set. Then the max-closure of the set \mathbf{Tree}_X of tree metrics on X is the set \mathbf{Met}_X of metrics on X .*

Proof. Since \mathbf{Met}_X is max-closed, it suffices to show that every $d \in \mathbf{Met}_X$ and every pair $xy \in \binom{X}{2}$ admit a metric $d' \in \mathcal{T}_X$ such that $d' \leq d$ and $d'_{xy} = d_{xy}$.

Let ρ denote Isbell’s isometric embedding of (X, d) in $\ell^\infty(X)$ [28], where for all $z \in X$ we have $\rho: z \mapsto \rho_z$ and $\rho_z(w) = d(z, w)$. Let $e_x: \ell^\infty(X) \rightarrow \mathbb{R}$ denote the evaluation map $e_x(f) = f(x)$ and set

$$d'_{zw} := |\rho_z(x) - \rho_w(x)| = |(e_x \circ \rho)(z) - (e_x \circ \rho)(w)|.$$

Thus, d' is a pull-back of the standard metric on the real line, hence a tree (semi-) metric. Since e_x is non-expansive and ρ is an isometry we also have $d' \leq d$. Finally,

at the same time we have

$$d'_{xy} = |\rho_x(x) - \rho_x(y)| = |0 - d_{xy}| = d_{xy},$$

which finishes the proof. \square

An immediate corollary is the non-existence of clustering projections — and hence, equivalently, of stationary sieving methods — whose image coincides (for all underlying sets X) with the set of tree metrics:

Corollary 38. *If \mathbf{D} is a clustering domain satisfying $\mathbf{D}_X \supseteq \mathbf{Tree}_X$ for all X , then $\mathbf{D}_X \supseteq \mathbf{Met}_X$ for all X .* \square

Thus, from the point of view of functorial clustering (if based on the metric category), the class of tree metrics — the first most accepted (and used) generalization of the class of dendrograms — is not suitable for characterizing “clustered” objects. Moreover, there exists no “middle ground” class of metrics lying between tree metrics and general metrics which could be used for this purpose. In particular, since $\mathbf{Tree}_X \subseteq \mathbf{Cut}_X \subseteq \mathbf{Met}_X$, there is no hope of obtaining a sieving method characterized by the class of cut metrics:

Corollary 39. *If \mathbf{D} is a clustering domain satisfying $\mathbf{D}_X \supseteq \mathbf{Cut}_X$ for all X , then $\mathbf{D}_X \supseteq \mathbf{Met}_X$ for all X .* \square

We now move on to examples of subcategories which are in fact clustering domains. In Section 3.2, we developed an explicit description of a projection functor for single-linkage. Although that projection highlights the naturality of the projection, there is also a computationally efficient algorithm for producing the maximal ultrametric under a given weight: the SLINK algorithm of Sibson [34]. On the other hand, the minimum spanning tree approach to single-linkage clustering in [24] provides an efficient way of producing clusters (*i.e.*, the \mathcal{SL} sieving functor). Additionally, the resulting sieves are in this case actually dendrograms, where the sieve itself carries important geometric features (see section Section 4 below). This leads us to investigate examples of clustering domains for which:

- (i) there is an efficiently computable projection,
- (ii) the associated sieving functor (after composing the projection with the Rips functor) is computationally tractable,
- (iii) the associated subcategory of **Sieve** has objects with a suitable geometric characterization.

In this section, we will identify several examples, focusing on the first two items, while the goal of Section 4 is to explore the possibility of extending the third.

Example 40 (Metric spaces). The category **Met** itself can be considered a clustering domain in **Weight**. The canonical projection $\mathcal{P}_{\mathbf{Met}}$ in this case is obtained by replacing the weights by the distance given by the path metric, *i.e.* the length of the shortest path between the points. This projection is then a well-studied algorithm with several standard computationally efficient approaches. Perhaps most importantly, this example of a clustering domain gives us a way to produce metrics by intersecting with other clustering domains. Care must be taken, however, as in general the projection to an intersection is not as simple as applying one projection and then the other. The associated sieving functor is not known to have an efficient

implementation, as it requires computing the maximal simplices of the Rips complex (or equivalently, the maximal cliques of the Rips graph) for arbitrary metric spaces.

Example 41 (Inframetrics). A weight space (X, u) will be called a ρ -inframetric space if

$$u_{xz} \leq \rho \max\{u_{xy}, u_{yz}\}$$

for every x, y, z in X .

Note that when $\rho = 1$, we recover the ultrametric condition, and that all metric spaces are 2-inframetric spaces (but some 2-inframetrics are not metrics). It is easy to see that the category of (finite) ρ -inframetric spaces is sup closed and forms a clustering domain.

For $1 < \rho < 2$, we can intersect the ρ -inframetric spaces with **Met** to obtain a clustering domain which contains inframetrics inside the category of all metric spaces. This would yield a projection functor whose associated sieving functor gives a functorial (overlapping) clustering method refining single-linkage. The authors are not aware of tractable algorithms for computing the general ρ -inframetric projections or the associated sieving functors. Moreover, the geometry of inframetric sieves (those sieves in the image of the associated sieving functor) does not appear to have a known characterization.

In a related vein, we could also work with the ρ -relaxed triangle inequality

$$u_{xz} \leq \rho(u_{xy} + u_{yz}).$$

The ρ -inframetric inequality implies the ρ -relaxed triangle inequality which in turn implies the 2ρ -inframetric inequality. The ρ -relaxed inequality also leads to a valid clustering domain.

Example 42 (q -Metric Spaces). Another example of a clustering domain is the subcategory **Met** $_q$ of q -metric spaces of Ribeiro, Segarra, Mémoli, and Carlsson defined in [33]. For $1 \leq q \leq \infty$, a q -metric space is a weight space (X, u) where u satisfies

$$(u_{xz})^q \leq (u_{xy})^q + (u_{yz})^q \text{ for } q < \infty$$

or

$$u_{xz} \leq \max\{u_{xy}, u_{yz}\} \text{ for } q = \infty.$$

When $q = 1$ we get ordinary metric spaces and when $q = \infty$ we get ultrametrics. Note that when $q = 2$, we have all triangles being acute. Recall for ultrametrics all triangles are isosceles with the longest side repeating. Again it is easy to see that q -metric spaces form a clustering domain, and a primary objective of [33] is to investigate the projection for $q < \infty$ defined by

$$(\mathcal{P}_{\mathbf{Met}_q} u)_{xy} = \min_{p \in \mathbb{P}(x, y)} \|p\|_q,$$

where $u \in \mathbf{Weight}$, $\mathbb{P}(x, y)$ is the set of edge-paths from x to y in the weighted distance graph, considered as vectors in \mathbb{R}^k (with k the length of the path) and $\|\cdot\|_q$ the usual q -norm. An efficient algorithm for the sieving functor and a geometric understanding of q -metric sieves are still needed.

Example 43 (Discretizations). Let **Int** be the category of finite weight spaces with integer distances. This is obviously sup closed in the sense that given any finite set X with a weight u there exists a unique maximal integer weight under u . When

restricting to metrics (*i.e.*, intersecting $\mathbf{Int} \cap \mathbf{Met}$, we must be careful to first take the integer floor of the distances given by u and then take the resulting path metric on X with those integer weights. The latter step is necessary since the floor of u may not be a metric. Also note that reversing this procedure is not necessarily the same.

Notice that this example can be generalized in a straightforward way using any given sequence of non-negative real numbers, providing a mechanism for discretizing in a way appropriate for a given application that does not destroy the underlying theoretical framework for clustering. Generically, an efficient algorithm for the sieving functor is bounded in complexity by an efficient algorithm for Rips sieving, since combinatorially these will produce the same sieve.

Example 44 (Quotient spaces). Let X be a fixed set and suppose \sim is an equivalence relation on X with quotient Z . Let $\pi: X \rightarrow Z$ denote the quotient map. Then $\pi^*: \mathbf{Met}_Z \rightarrow \mathbf{Met}_X$ embeds \mathbf{Met}_Z in \mathbf{Met}_X as the set of all metrics satisfying $w_{xy} = 0$ whenever $x \sim y$ in X .

Observe that $\pi^*\mathbf{Met}_Z$ is a closed and max-closed subset of \mathbf{Met}_X , hence also sup-closed, by setting Proposition 18. Define a functor $\mathcal{F}_X: \mathbf{Met}_X \rightarrow \mathbf{Met}_X$ by $\mathcal{F}_X(w) = \sup\{u \in \pi^*\mathbf{Met}_Z \mid u \leq w\}$. Then $\mathcal{F}_X^2 = \mathcal{F}_X$, $\mathcal{F}_X(\mathbf{Met}_X) = \pi^*\mathbf{Met}_Z$, $\mathcal{F}_X w \leq w$, and $w_1 \leq w_2$ implies $\mathcal{F}_X w_1 \leq \mathcal{F}_X w_2$ for all $w, w_1, w_2 \in \mathbf{Met}_X$. Any map $\mathcal{H}: \mathbf{Met}_X \rightarrow \mathbf{Met}_X$ satisfying these conditions must coincide with \mathcal{F}_X .

For any clustering domain \mathbf{D} contained in \mathbf{Met} , we have the diagram:

$$\begin{array}{ccc}
 \mathbf{Met}_Z & \xleftarrow{\pi^{*-1} \circ \mathcal{F}_X} & \mathbf{Met}_X \\
 \downarrow (\mathcal{P}_{\mathbf{D}})_Z & \xrightarrow{\pi^*} & \downarrow (\mathcal{P}_{\mathbf{D}})_X \\
 \mathbf{Met}_Z & \xleftarrow{\pi^{*-1} \circ \mathcal{F}_X} & \mathbf{Met}_X
 \end{array}$$

A red curved arrow labeled \mathcal{G}_X connects the bottom \mathbf{Met}_X to the top \mathbf{Met}_X , passing through the right side of the diagram.

We claim that $\mathcal{G}_X = \pi^{*-1} \circ \mathcal{F}_X \circ (\mathcal{P}_{\mathbf{D}})_X \circ \pi^*$ coincides with the map $(\mathcal{P}_{\mathbf{D}})_Z$. Indeed, it satisfies all 3 requirements for characterizing $(\mathcal{P}_{\mathbf{D}})_Z$. Thus \mathcal{F}_X is a natural “quotient operation” on metric spaces that commutes with $\mathcal{P}_{\mathbf{D}}$.

Recall that, for any metric $w \in \mathbf{Met}_X$ the quotient metric \bar{w} on the quotient Z is defined as follows. First, for $A, B \in Z$ one considers ‘paths’ from A to B :

$$\mathbb{P}(A, B) := \left\{ (x_0, y_1, x_1, \dots, y_{n-1}, x_{n-1}, y_n) \mid \begin{array}{l} x_0 \in A, y_n \in B, n \in \mathbb{N}, \\ x_i \sim y_i \text{ for } i = 1, \dots, n \end{array} \right\}$$

Next, for each $p \in \mathbb{P}(A, B)$ as above one defines its length as $\ell(p) := \sum_{i=1}^n w(x_{i-1}, y_i)$. Finally:

$$\bar{w}_{AB} := \inf_{p \in \mathbb{P}(A, B)} \ell(p).$$

Since $\pi^*(\bar{w}) \leq w$ for all $w \in \mathbf{Met}_X$, $\mathcal{F}_X(w)$ coincides with $\pi^*(\bar{w})$. This example is a bit different from those above, but the interesting point here is that quotient metrics respect *any* stationary functorial sieving map, providing yet another reason to adopt this particular generalization of hierarchical clustering.

Example 45 (Any subcategory of **Weight**). We remind the reader that any subcategory of **W** with the pull-back property can be sup closed to give a clustering domain. The difficult issue is to understand exactly what that closure is and how to construct the projection to it in a practical way. With the ability to take intersections of clustering domains, it is clear that there are many such subcategories, providing different clustering methods. Much more work is needed to characterize additional useful clustering domains.

4. INJECTIVE ENVELOPES AND THE ROLE OF ANTIPODES IN CLUSTERING

The significance of Corollaries 38 and 39 is in their prohibiting a functorial compromise between the classical notion of hierarchical clustering and the geometric clustering maps based on split decompositions championed by the Dress school. However, the rather complicated relation between split-based clustering methods and the geometry of injective envelopes inspires hope that a sieving functor having some of the qualities of split-based maps could be derived from the geometry of the injective envelopes of spaces chosen to lie in an appropriately defined clustering domain. Setting aside all technical details for a separate account [12], we focus in this section on introducing and reviewing relevant properties of the clustering domain of *A-spaces*, motivated by a study of rooted injective envelopes (see below), as well as some additional clustering sub-domains which, we believe, will prove useful as a means of constructing new classes of hierarchical classifiers lying between the category of dendrograms provided by \mathcal{SL} and the category of unrestricted sieves provided by \mathcal{R} .

4.1. Injective envelopes and clustering. Recall¹ that a metric space X is said to be *injective*, if, for any isometry $i: A \rightarrow B$, any non-expansive map $f: A \rightarrow X$ extends to a non-expansive map $F: B \rightarrow X$ (in the sense that $F \circ i = f$). As any metric space X embeds isometrically in an injective one (consider the Banach embedding), one asks whether an *injective envelope* ϵX exists, that is: an isometry $e: X \rightarrow \epsilon X$ into an injective space ϵX , through which any embedding of X to an injective space must factor. The construction of ϵX , discovered independently by Isbell [28] and Dress [19], is explicit and easy to describe: first setting

$$P(X, d) := \left\{ f: X \rightarrow \mathbb{R}_{\geq 0} \mid \forall x, y \in X \ f(x) + f(y) \geq d_{xy} \right\}$$

one then lets ϵX be the subset of those $f \in P(X, d)$ that are pointwise minimal, inheriting its metric from the sup-metric. The mapping $e: X \rightarrow \epsilon X$ defined by $e(x)(y) = d_{xy}$ is the required embedding.

The injective envelope $e: X \rightarrow \epsilon X$ provides a canonical way to minimally “fill-in” X so that as many of its points as possible appear as endpoints of the complete and hyper-convex [1] metric space ϵX , prompting the idea that cluster hierarchies X might be derived from connectivity properties of ϵX . For example, computing the cut-point hierarchy [37] of ϵX [20] induces a natural nested tree-like structure on X in the sense of [15]. The study of the relation between envelopes and split decompositions by the Dress school [3, 16, 21, 17] following Bunemann’s work [7] pushes this idea even further, but seems to lose track of the projective/hierarchical aspects of the clustering problem.

¹For a modern, self-contained account of injective envelopes see [32].

4.2. Antipodes. Our present outlook on the usefulness of injective envelopes is motivated by two well known facts. First, any dendrogram (X, θ) — when realized as a metric tree with leaf set X , in which the length of each edge is equal to half the absolute value of the difference in heights between its endpoints — is naturally isometric to the injective envelope of (X, u) , where $u \in \mathbf{Ult}_X$ is the ultra-metric corresponding to the given dendrogram, $(X, \theta) = \mathcal{R}(X, u)$. Second, a metric tree is the geometric realization of a dendrogram if and only if it contains a *root vertex*: a point at equal distances from all the points of X (also known as the *leaves* of the dendrogram). We are led to investigate (compact metric) spaces with the property that their injective envelope has a unique root in the same sense. It turns out that these are precisely the (compact, metric) spaces where every point has an antipode, or *A-spaces*:

Definition 46 (Antipode, *A-space*²). A metric space (X, d) is said to be an *A-space* if every $x \in X$ has an *antipode*, that is: there exists $y \in X$ with $d_{xy} = \text{diam}(X, d)$. Denote the set of *A-spaces* with underlying set X by \mathbf{A}_X .

Remark 47. For finite X , the collection \mathbf{A}_X is a closed subspace of \mathbf{Weight}_X , and is clearly max-closed, hence also sup-closed. Note that $(X, d) \in \mathbf{Ult}$ if and only if every subspace of (X, d) is an *A-space*, or equivalently, if every triple in (X, d) is an *A-space*. In particular, an ultrametric space is globally an *A-space*.

Overall, we conclude that the class of *A-spaces* forms a clustering domain fibering over the the category of sets and *surjective* set maps. The surjectivity restriction is necessary, because only for surjective set maps is it true that the pull-back of an *A-metric* is again an *A-metric*. The surjectivity restriction is also more than a mere convenience to accommodate the category-theoretical lingo. In fact, it results in relaxing the consistency requirements on the corresponding sieving functor from **Met** to **Sieve**, implying that the resulting clustering method is consistent with respect to quotients, but not necessarily consistent with respect to sub-sampling (which corresponds to pull-backs via injective maps).

The property of X being an *A-space* is, in fact, encoded in the geometry of its injective envelope:

Theorem 48. *Given a compact metric space (X, d) , let \hat{X} denote the extension $X \cup \{\infty\}$, $\infty \notin X$, adding $d_{\infty x} := \frac{1}{2}\text{diam}(X, d)$ for all $x \in X$. The following are equivalent:*

- (1) (X, d) is an *A-space*;
- (2) $\epsilon\hat{X}$ is isometric to ϵX through an isometry fixing X pointwise;
- (3) ϵX contains a point — denoted ∞_X — at equal distances to all points of X ;
- (4) ϵX contains a point at a distance $\frac{1}{2}\text{diam}(X, d)$ to every $x \in X$.

The point ∞_X will be referred to as the root of (X, d) in ϵX . □

The example of a dendrogram viewed as the injective envelope of an ultra-metric space (X, d) provides a hint at the class of subspaces one might want to regard as clusters for the general *A-space*. When ϵX is a dendrogram, each $f \in \epsilon X$ is no more than a specification of distances to the leaves, with the leaves closest to f — and hence of minimal value under f — forming the associated ‘descendant’ cluster. See Figure 1 for a comparison in four-point spaces.

²The notion “Antipodal Space” has already been put to extensive good use in [27, 26]




FIGURE 1. Envelopes of the two generic types of 4-point A -spaces (a,b), seen as ‘fatted’ versions of a dendrogram (c), also drawn in a more standard representation in (d). The root of each envelope is marked with a red \otimes . For a more meaningful comparison, the metric in (a-c) is chosen to be the path metric induced from embedding in the ℓ_1 plane.

Denote the minimum value of f by $\|f\|$, and the set of points where $\|f\|$ is achieved by $\text{Min}(f)$. For a general A -space we have:

Proposition 49. *Suppose (X, d) is an A -space, then every $f \in \epsilon X$ satisfies*

$$\text{dist}(\infty_X, f) = \frac{1}{2} \text{diam}(X, d) - \|f\|$$

Moreover, every $f \in \varepsilon(X, d)$ satisfies

$$\max_{z \in X} f(z) + \|f\| = D$$

In particular, $f(x) + f(y) = D$ whenever $x \in \text{Min}(f)$ and y is an antipode of x . \square

Thus, A -spaces form a clustering domain containing the ultrametric clustering domain and generalizing some pertinent geometric properties of ultrametrics. Moreover:

Remark 50. Given a metric space (X, d) , its projection to the category of A -spaces may be computed recursively as follows. Set $d_0 = d$ and for any $t \geq 0$ define $E_t = \{xy \mid d_t(x, y) = \text{diam}(X, d_t)\}$. If E_t is an edge cover then stop and return d_t , as (X, d_t) is an A -space if and only if E_t is an edge cover of X . Else, for each $xy \in E_t$ set $d_{t+1}(x, y)$ to equal the second-largest distance in (X, d_t) ; for $xy \notin E_t$ set $d_{t+1}(x, y) = d_t(x, y)$.

4.3. Clustering domains of A -spaces.

4.3.1. *4-point conditions, and more.* Additional classes of A -spaces exist, forming clustering domains which contain the domain of ultra-metrics.

Definition 51. Let $m \geq 3$ be an integer. We say that a metric space (X, d) is an A^m -space, if every subset of cardinality m is an A -space.

The following lemma is easy to prove:

Lemma 52. *Let $m \geq 3$ be an integer. If (X, d) is an A^m -space, then it also an A^n -space for any $n \geq m$. In particular, a finite A^m -space of cardinality at least m is an A -space.* \square

It is also straightforward to observe that pull-backs of A^n -spaces under injective maps remain A^n -spaces, as well as that the set of A^n -spaces with a fixed finite base space X is a closed and max-closed subset of \mathbf{Weight}_X . Therefore:

Corollary 53. *For every integer $m \geq 3$, the class of finite metric spaces that are A^m -spaces forms a clustering domain over the category of sets and injective maps. This clustering domain will be denoted by \mathbf{A}^m .* \square

The resulting tower of clustering domains over the category of sets and *injective* set maps,

$$\mathbf{A}^3 \subsetneq \mathbf{A}^4 \subsetneq \cdots \subsetneq \cdots \mathbf{A}^m \cdots,$$

is bounded above by the *class* \mathbf{A} of compact A -spaces, though we must be careful not to view \mathbf{A} as a clustering domain in this context, since the morphism structures of \mathbf{A}^m and \mathbf{A} cannot be reconciled. Similarly, while $\mathbf{A}^3 = \mathbf{Ult}$ as a class of objects (see Remark 47), the inclusion map from \mathbf{Ult} to \mathbf{A}^4 is not a functor, as \mathbf{Ult} admits all non-expansive set maps—not just the injective ones. Because the poset structure of the fibers are consistent (see Remark 5), however, we are still able to consider the factorization of the projections $\mathcal{P}_{\mathbf{A}^m}$ through $\mathcal{P}_{\mathbf{A}}$.

Thus, the sieving functors from \mathbf{A}^m to \mathbf{Sieve} should be viewed as functors which are only consistent with respect to sub-sampling. We have yet to obtain a projection algorithm to, say, \mathbf{A}^4 -spaces, or a geometric characterization of their injective envelopes.

4.4. Discussion: Injective envelopes and practical clustering. The view of a sieving functor as a projection in the weight/metric category followed by Rips clustering is attractive from the point of view of it providing a relatively simple and geometrically motivated means for constructing such functors *abstractly*, but from a computational standpoint it is a bit naïve. Indeed, the computational complexity of Rips sieving is, essentially, prohibitive for large data sets. One therefore is after clustering domains whose geometric properties enable efficient algorithms for computing the clustering directly.

In this context, the domain of ultra-metrics provides an extreme example, where computing the sieve associated with a given metric space is not just efficient, but can be efficiently distributed/decentralized.

On the other extreme, the domain of A -spaces seems to offer little improvement over Rips sieving of a general metric: every metric space (X, d) extends to an A -space by adding just one more point; in other words, though geometrically well-motivated, the requirement that (X, d) be an A -space is not sufficiently restrictive. The algorithm in Remark 50 suggests that unless the projection (X, d_A) of (X, d) to the domain of A -spaces yields a clear and roughly even *partitioning* of X into clusters at some degree of resolution, Rips-clustering of (X, d_A) will still require a search through a significant portion of the 2^n possible clusters.

Clearly, this ‘malfunction’ of sieving through A -spaces is due to a lack of some kind of hereditary/hierarchical structure in spaces in this category: no particular collection of subspaces of an A -space is forced to inherit the A -space condition. This motivates the introduction of the notion of A^m -spaces, best viewed as a relaxation of the hereditary class of ultra-metric spaces — A^4 -spaces being of particular interest, in view of the key role of 4-point conditions in the metric clustering literature [7, 19].

Judging from the ultra-metric case, it seems plausible that, for a clustering domain \mathbf{D} , the existence of an efficient sieving algorithm requires a combination of (1) geometric constraints forcing some notion of ‘thinness’ on $\epsilon(X, d)$ whenever $(X, d) \in \mathbf{D}$; (2) all clusters of $(X, d) \in \mathbf{D}$ lying in \mathbf{D} , similarly to the notion of *excisiveness* introduced in [9]; and (3) proper restrictions on the base morphisms.

The extent to which this vague conjecture holds true for \mathbf{A}^m , $m \geq 4$, is the subject of ongoing work.

REFERENCES

- [1] Nachman Aronszajn and Prom Panitchpakdi. Extension of uniformly continuous transformations and hyperconvex metric spaces. *Pacific J. Math*, 6(3):405–439, 1956.
- [2] Hans-Jürgen Bandelt and Andreas W. M. Dress. Weak hierarchies associated with similarity measures: An additive clustering technique. *Bull. Math. Biol.*, 51(1):133–166, 1989.
- [3] Hans-Jürgen Bandelt and Andreas W. M. Dress. A canonical decomposition theory for metrics on a finite set. *Adv. Math.*, 92:47–105, 1992.
- [4] Hans-Jürgen Bandelt and Andreas W. M. Dress. An order theoretic framework for overlapping clustering. *Discrete Math.*, 136(1–3):21–37, December 1994.
- [5] Patrice Bertrand. Set systems and dissimilarities. *European J. Combin.*, 21:727–743, 2000.
- [6] Patrice Bertrand and Jean Diatta. Weak hierarchies: A central clustering structure. In Fuad Aleskerov, Boris Goldengorin, and Panos M. Pardalos, editors, *Clusters, Orders, and Trees: Methods and Applications*, Springer Optimization and Its Applications, pages 211–230. Springer New York, 2014.
- [7] O. Peter Buneman. The recovery of trees from measures of dissimilarity. In D. G. Kendall and P. Tautu, editors, *Mathematics in the Archaeological and Historical Sciences*. Edinburgh University Press, 1971.
- [8] Gunnar Carlsson and Facundo Mémoli. Characterization, stability, and convergence of hierarchical clustering methods. *J. Mach. Learn. Res.*, 11:1425–1470, 2010.
- [9] Gunnar Carlsson and Facundo Mémoli. Classifying clustering schemes. *Found. Comput. Math.*, 13(1):221–252, 2013.
- [10] Frédéric Chazal, Steve Oudot, Primoz Skraba, and Leonidas J. Guibas. Persistence-based clustering in Riemannian manifolds. *J. ACM*, 60(6), 2013.
- [11] Jared Culbertson, Dan P. Guralnik, Jakob Hansen, and Peter F. Stiller. Consistency constraints for overlapping data clustering. *arxiv preprint arXiv:1608.04331*, 2016.
- [12] Jared Culbertson, Dan P. Guralnik, and Peter F. Stiller. Antipodal spaces, injective envelopes and clustering. *in preparation*, 2016.
- [13] Michel Marie Deza and Monique Laurent. *Geometry of cuts and metrics*, volume 15 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1997.
- [14] Jean Diatta. One-to-one correspondence between indexed cluster structures and weakly indexed closed cluster structures. In Paula Brito, Patrice Bertrand, Guy Cucumel, and Francisco de Carvalho, editors, *Selected Contributions in Data Analysis and Classification*, Studies in Classification, Data Analysis, and Knowledge Organization, pages 477–482. Springer Berlin Heidelberg, 2007.
- [15] Warren Dicks and Martin John Dunwoody. *Groups acting on graphs*, volume 17. Cambridge University Press, 1989.
- [16] A. Dress, K. Huber, and V. Moulton. Some variations on a theme by Buneman. *Ann. Comb.*, 1:339–352, 1997.
- [17] Andreas Dress, Vincent Moulton, Andreas Spillner, and Taoyang Wu. Obtaining splits from cut sets of tight spans. *Discrete Appl. Math.*, 161:1409–1420, 2013.
- [18] Andreas W. M. Dress, Katarina T. Huber, and Vincent Moulton. Totally split-decomposable metrics of combinatorial dimension two. *Ann. Comb.*, 5(1):99–112, 2001.
- [19] Andreas WM Dress. Trees, tight extensions of metric spaces, and the cohomological dimension of certain groups: a note on combinatorial properties of metric spaces. *Adv. Math.*, 53(3):321–402, 1984.
- [20] Andreas WM Dress, Katharina T Huber, Jacobus Koolen, and Vincent Moulton. An algorithm for computing virtual cut points in finite metric spaces. In *International Conference on Combinatorial Optimization and Applications*, pages 4–10. Springer, 2007.
- [21] Andreas WM Dress, KT Huber, and Vincent Moulton. A comparison between two distinct continuous models in projective cluster theory: The median and the tight-span construction. *Ann. Comb.*, 2(4):299–311, 1998.
- [22] Brian Everitt. *Cluster Analysis*. Wiley series in probability and statistics. Wiley, 5 edition, 2011.

- [23] Fernando Gama, Santiago Segarra, and Alejandro Ribeiro. Overlapping clustering of network data using cut metrics. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015.
- [24] John C. Gower and G. J. S. Ross. Minimum spanning trees and single linkage cluster analysis. *J. R. Stat. Soc. Ser. C. Appl. Stat.*, pages 54–64, 1969.
- [25] W Harvey, O Rübel, V Pascucci, P-T Bremer, and Y Wang. Enhanced topology-sensitive clustering by Reeb graph shattering. In Ronald Peikert, Helwig Hauser, Hamish Carr, and Raphael Fuchs, editors, *Topological Methods in Data Analysis and Visualization II: Theory, Algorithms, and Applications*, Mathematics and Visualization, pages 77–90. Springer Berlin Heidelberg, 2012.
- [26] Katharina T Huber, Jacobus H Koolen, and Vincent Moulton. The tight span of an antipodal metric space: part II - geometrical properties. *Discrete Comput. Geom.*, 31(4):567–586, 2004.
- [27] Katharina T Huber, Jacobus H Koolen, and Vincent Moulton. The tight span of an antipodal metric space: part I - combinatorial properties. *Discrete Math.*, 303(1):65–79, 2005.
- [28] J. R. Isbell. Six theorems about injective metric spaces. *Comment. Math. Helv.*, 39(1):65–76, 1964.
- [29] Melvin F. Janowitz. *Ordinal and Relational Clustering*, volume 10 of *Interdisciplinary Mathematical Sciences*. World Scientific, 2010.
- [30] Nicholas Jardine and Robin Sibson. *Mathematical Taxonomy*. Wiley Series in Probability and Mathematical Statistics: Tracts on Probability and Statistics. Wiley, 1971.
- [31] Jon Kleinberg. An impossibility theorem for clustering. *Advances in Neural Information Processing Systems*, 15, 2002.
- [32] Urs Lang. Injective hulls of certain discrete metric spaces and groups. *J. Topol. Anal.*, 05(03):297–331, 2013.
- [33] Santiago Segarra, Gunnar Carlsson, Facundo Mémoli, and Alejandro Ribeiro. Metric representations of network data. *preprint*, 2015.
- [34] Robin Sibson. SLINK: an optimally efficient algorithm for the single-link cluster method. *The Computer Journal*, 16(1):30–34, 1973.
- [35] Gurjeet Singh, Facundo Mémoli, and Gunnar Carlsson. Topological methods for the analysis of high dimensional data sets and 3D object recognition. In *Symposium on Point-Based Graphics*, pages 91–100, 2007.
- [36] Bernd Sturmfels and Josephine Yu. Classification of six-point metrics. *Electron. J. Combin.*, 11, 2004.
- [37] L. E. Ward. Axioms for cutpoints. In *General topology and modern analysis (Proc. Conf., Univ. California, Riverside, Calif., 1980)*, pages 327–336, New York-London, 1981. Academic Press.
- [38] Dekang Zhu, Dan P Guralnik, Xuezhi Wang, Xiang Li, and Bill Moran. Statistical properties of the single linkage hierarchical clustering estimator. *arXiv preprint arXiv:1511.07715*, 2015.

SENSORS DIRECTORATE, AIR FORCE RESEARCH LABORATORY
E-mail address: `jared.culbertson@us.af.mil`

SCHOOL OF ENGINEERING AND APPLIED SCIENCE, UNIVERSITY OF PENNSYLVANIA
E-mail address: `guraldan@seas.upenn.edu`

DEPARTMENT OF MATHEMATICS, TEXAS A&M UNIVERSITY
E-mail address: `stiller@math.tamu.edu`